All that begins . . .

السلام عليكم

peace be upon you

UTM
UNIVERSITI TEKNOLOGI MALAYSIA

SCAE.Initiative
(Open Source Computer Aided Engineering)

Engineering Computing Linux Cluster

Parallel Computing – An Overview

Abu Hasan 'ABDULLAH

December 2019

Outline

Computer Cluster

- A computer cluster is a group of linked computers, working together closely so that in many respects they form a single computer.
- The components of a cluster are commonly, but not always, connected to each other through fast local area networks.
- Clusters are usually deployed to improve performance and/or availability over that of a single computer.

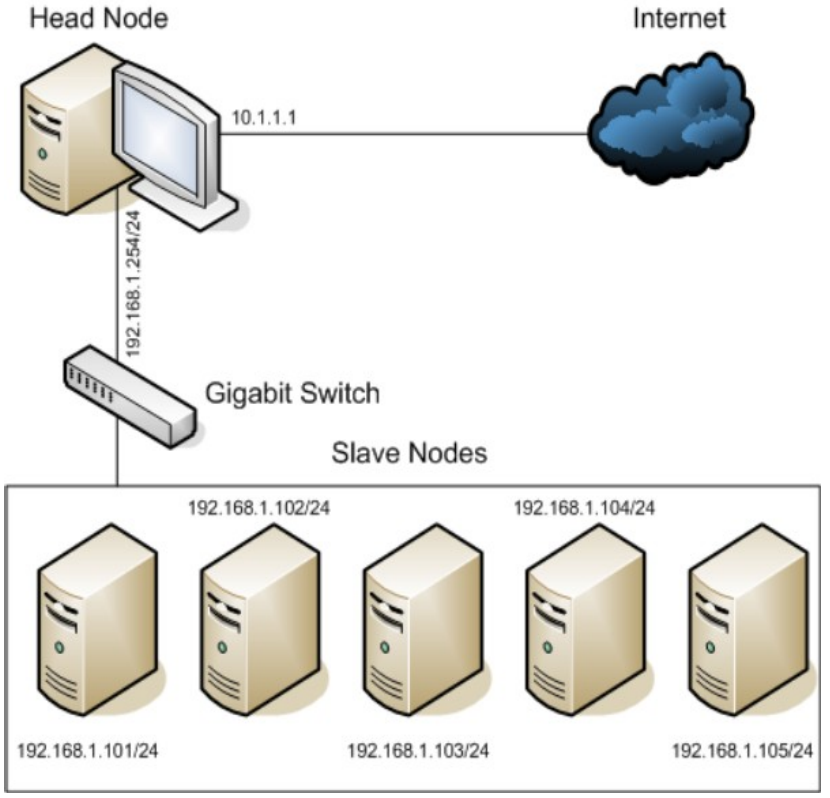

http://en.wikipedia.org/wiki/Cluster_(computing)

What is Parallel Computing?

- Traditional **serial** computation
  - Run on a single computer having a single Central Processing Unit (CPU).
  - A problem is broken into a discrete series of instructions.
  - Instructions are executed one after another.
  - Only one instruction may execute at any moment in time.
- **Parallel** computing is the simultaneous use of multiple compute resources to solve a computational problem:
  - To be run using multiple CPUs or multicore CPUs
  - A problem is broken into discrete parts that can be solved concurrently
  - Each part is further broken down to a series of instructions
  - Instructions from each part execute simultaneously on different CPUs

What is Parallel Computing?

- **Parallel** computing resources can include:
  - A single computer with multiple processors;
  - An arbitrary number of computers connected by a network;
  - A combination of both.

Uses for Parallel Computing

- Atmosphere, Environment, Earth
- Physics—applied, nuclear, particle, condensed matter, high pressure, fusion, photonics
- Bioscience, Biotechnology, Genetics
- Chemistry, Molecular Sciences
- Geology, Seismology
- Mechanical Engineering—from prosthetics to spacecraft
- Electrical Engineering, Circuit Design, Microelectronics
- Computer Science, Mathematics

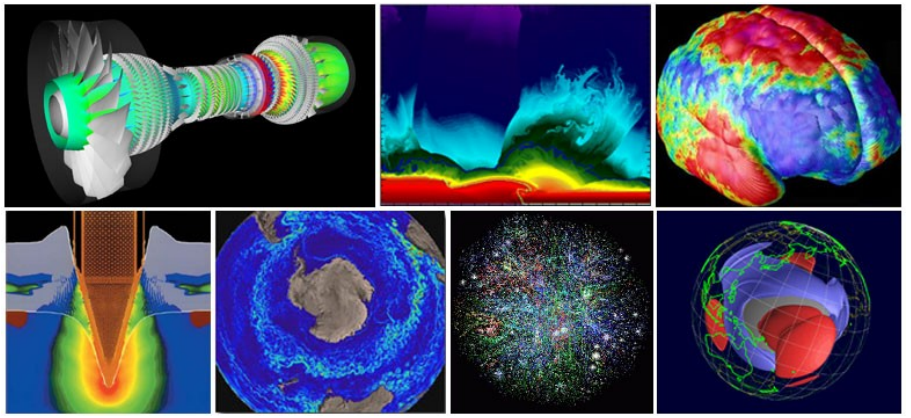

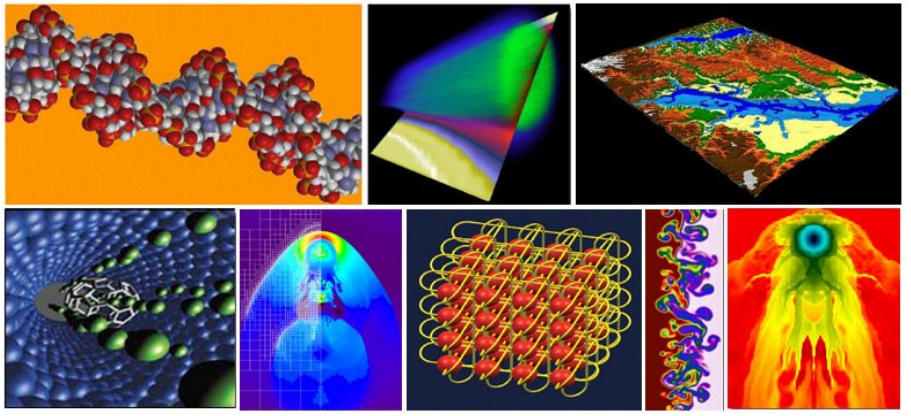

Why Parallel Computing?

- Save time and/or money—more resources at a task shortens time to complete, with potential cost savings. Parallel clusters can be built from cheap, commodity components.
- Solve larger problems—large and/or complex that it is impractical or impossible to solve them on a single computer, especially given limited computer memory.
- Provide concurrency—a single compute resource can only do one thing at a time. Multiple computing resources can be doing many things simultaneously.
- Use of non-local resources—using compute resources on a wide area network, or even the Internet when local compute resources are scarce.
- Limits to serial computing—physical and practical reasons pose significant constraints to simply building ever faster serial computers

Parallel Programming Models

- Parallel programming models in common use:
  - Shared Memory
  - Threads
  - Message Passing
  - Data Parallel
  - Hybrid
- There is no “best” model, although there certainly are better implementations of some models over others.
- Our main interest is on the Message Passing model on a shared memory machine

Parallel Programming Models

- Characteristics of Message Passing model:
  - A set of tasks that use their own local memory during computation. Multiple tasks can reside on the same physical machine as well across an arbitrary number of machines.
  - Tasks exchange data through communications by sending and receiving messages.
  - Data transfer usually requires cooperative operations to be performed by each process. For example, a send operation must have a matching receive operation.
- Implementations:
  - Message Passing Interface (MPI) Forum was formed (1992) to establish a standard interface for message passing implementations.
  - Part 1 of the MPI was released in 1994.
  - Part 2 (MPI-2) was released in 1996.

Tools to Build a Parallel Cluster

- Multiuser, Multitasking Operating System
  - Linux, Unix
  - Windows Server
- MPI Implementations
  - OpenMPI
  - LAM-MPI
- Compilers
  - GNU C/C++/Fortran
  - Intel C/C++/Fortran
  - PGI C/C++/Fortran
- Schedulers
  - OpenPBS
  - Torque
  - Sun Grid Engine (SGE)
- Monitor(s)
  - Ganglia
  - Munin

CAE Tools for Linux Cluster

Commercial Packages

- Ansys Fluent
- Matlab
- Abaqus
- Maple/Mathematica
- COMSOL Multiphysics
- Intel C/C++/Fortran
- Pointwise/Gambit/Amira
- EnSight/Tecplot
- Maya
- SPSS

OSS Packages

- Blender/FreeCAD
- gmsh/EnGrid/MeshLab/Salome
- Code_Saturne/OpenFOAM
- Code_Aster/Elmer/z88
- FreeShip/NetGen/NGSolve
- ParaView/VisIT/GNUplot
- Octave/Scilab/Maxima/SciDAVis
- GNU C/C++/Fortran/Python
- R/PSPP

. . . must end

- . . . and I end my presentation with two supplications

رَّبِّ زِدْنِي عِلْمًا

my Lord! increase me in knowledge

(TAA-HAA (20):114)

اَللّٰهُمَّ اِنَّا نَسْئَلُكَ عِلْمًا نَافِعًا

O Allah! We ask You for knowledge that is of benefit

(IBN MAJAH)